Institute for Scientific Computing Research

Workshop and Conference Reports

Synopsis of Workshop Events:

Bay Area Scientific Computing Day

The Bay Area Scientific Computing Day is an annual gathering to encourage the interaction and collaboration of researchers in the field of scientific computing from the Bay Area. It also provides researchers who are new to the Bay Area an opportunity to present their work to the local community. BASCD 2001 was held on Saturday, February 10, and was hosted by Lawrence Livermore National Laboratory and organized by CASC and the ISCR. More than 120 researchers attended the event, from laboratories, universities, and industry. Thirteen technical presentations were given by graduate students, postdoctoral researchers, and other scientists nominated or invited from throughout the Bay Area.

Program details are available at the BASCD 2001 website located at: http://www.llnl.gov/CASC/workshops/bascd/

Synopsis of Workshop Events:

Workshop on Object-Oriented and Component Technology for Scientific Computing

The first Workshop on Object-Oriented and Component Technology for Scientific Computing was held July 23-25 in Livermore, California. Forty-six active researchers and developers of scientific software discussed the state of the art in the field of high-performance software technology. Participants came from a variety of backgrounds, including thirty-one from the DOE, three from NASA, eight from universities, and four from industry.

This workshop was divided into five sessions focusing on software issues that arise in scientific computing: scripting, component technology, distributed computing, problem solving environments (PSEs), and applications and libraries. Scripting provides a flexible and interactive mechanism to couple and steer large science simulations. Component technology provides a means to manage the complexity of modern scientific simulation software and enables new simulation capabilities that were previously unavailable due to limitations in interoperability. Distributed computing frameworks provide the illusion of homogeneity in a heterogeneous environment. Problem-solving environments provide rich functionality for development and deployment of domain specific software. Applications and libraries are the final products (and from the point of view of component technology developers) for the customers of these technologies. Each of these five sessions consisted of four or five speakers. A one-hour moderated panel discussion covering the general topic of the session followed.

This format worked exceedingly well because questions for the panel were submitted in advance when participants registered. These “killer questions” were advertised on the workshop’s website, and the speakers were reminded that if they did not address these questions in their talks, they would almost certainly have to do so in the panel discussions that followed.

Among the attendees, there quickly grew up a common theme about the difficulty of software development and integration. Very few speakers discussed their software in isolation. Most talks highlighted either how their software could be integrated into other packages, or how they used other packages to augment their own software.

It became clear that software reuse is a critical necessity for scientific computing to continue to advance. To meet the push for more physics, higher fidelity, and more functionality software has to become more reusable… and usable in broader contexts.

Workshop participants identified three major barriers to effective software reuse: interoperability, semantics, and portability.

- Software interoperability is hindered by language incompatibilities, the lack of standard scientific interfaces, failure to adhere to existing standards, software build and installation incompatibilities, and too strong assumptions much about the local computational environment.

Synopsis of Workshop Events (continued):

- Participants agreed that not enough practical research was being done in the area of scientific software semantics. Most current compilers and source-to-source translaters specify only the calling syntax of a scientific library. Semantics descriptions would extend that syntactic description to specify the behavior and valid usage of a software library.

- Finally, participants identified portability as a difficult and time-consuming task. Portability involves building software and the associated compilation system that is robust enough to work on multiple computer platforms. Portability is especially a concern in high-performance computing, since parallel platforms have a much shorter lifespan than the software, and the operating systems for most parallel platforms are specifically modified for the parallel programming environment.

The timing of this workshop was auspicious, as it came on the heels of the SciDAC awards and provided a showcase for CASC's role in developing Component Technology for scientific computing.

Program details are located at the conference website at http://www.llnl.gov/CASC/workshops/components_2001/

Synopsis of Workshop Events:

14th Copper Mountain Conference on Multigrid Methods

The 14th Copper Mountain Conference on Multigrid Methods was held on April 1-6, 2001. Over one hundred applied mathematicians and computer scientists from eleven countries attended the meeting and over thirty graduate students and post-docs were supported. During the five-day meeting, 68 talks on current research topics were presented, as well as three tutorials preceding the reception on Sunday.

Presentations with similar content were organized into sessions, including:

- Transport
- Inverse problems
- Algebraic methods
- Domain decomposition
- Nonlinear methods
- Fluid dynamics
- Ocean dynamics
- Multilevel theory
- High-performance computational strategies
- First-order systems of least squares

In addition, a lively Circus (of extemporaneous presentations, in the tradition of the Finite Element Circus) was convened Thursday night. All of the sessions, including the Circus and the tutorials, were very well attended. One of the hallmarks of the Copper conference series is the relaxed atmosphere and open, active discussions that it fosters. This collaborative environment was again very much in evidence at this year's meeting.

A student paper competition stimulated student participation in the Conference. This year, there were 25 students in attendance and about 15 returning scientists who first attended the Copper conference series as students. For the competition, students were asked to submit a singly authored paper containing original research. A panel of judges made up of members of the Program Committee selected four winners. These winners were: Avraham Kenigsberg, Malik Silva, Chisup Kim, and Markus Korwarschik.

Selected papers from the conference will appear following a standard review process in E.T.N.A.

Program details are located at the conference website at http://amath.colorado.edu/faculty/copper/2001/

Synopsis of Workshop Events:

The Third Workshop on Mining Scientific Datasets

The Third Workshop on Mining Scientific Datasets was held April 7, 2001, in conjunction with the First SIAM International Conference on Data Mining. The goal of the workshop was to bring together researchers from the data mining community and various science and engineering communities in order to better understand how data mining can be used for the exploration of scientific datasets. In particular, we hoped to identify the common threads across the diverse technical areas that could be effectively harnessed to solve the problems of scientific data analysis.

Unlike the previous two workshops in the series, where many of the talks were given by invited speakers, we solicited short papers on relevant subjects for the third workshop. These were reviewed by the workshop organizing and program committees. Of the four workshops that were held on the final day of the conference, the scientific data mining one had the largest number of registered attendees. Though the workshop was held on a Saturday, most of the 75 attendees remained throughout.

Each of the morning and afternoon sessions started with a keynote talk. The first keynote talk was by Dennis DeCoste from JPL, who described the work being done at NASA on "Mining Large Datasets using Support Vector Machines." The second keynote talk of the workshop was given by Sara Graves from the University of Alabama at Huntsville on "Creating an Environment for Scientific Data Mining." She described the ADAM toolkit that has been developed at UAH for remote sensing data.

These keynote talks complemented the contributed talks very well. There were 14 contributed talks, which covered a diverse set of subjects including mining astronomical data, multiresolution representation of structural mechanics data, use of boosting techniques to classify remotely sensed data, archiving and processing of remotely sensed data, and the mining of corn-field data to detect insurance fraud. Several of the talks were presented by students on their ongoing research work. During the breaks, it was interesting to see the astronomers exchanging ideas with the remote sensing data analysts about the use of a common set of tools for both "look up" and "look down" data.

The workshop was supported by the Army High Performance Computing Research Center and the Center for Applied Scientific Computing at the Lawrence Livermore National Laboratory. The latter provided funds for travel for eleven students and the keynote speakers. The co-organizers were Michael Burl (NASA Jet Propulsion Laboratory), Chandrika Kamath (Lawrence Livermore Laboratory), Vipin Kumar (University of Minnesota and AHPCRC), and Raju Namburu (Army Research Laboratory).

A web page with details on the workshop is at http://www.ahpcrc.umn.edu/conferences/

Synopsis of Workshop Events:

The 2001 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Industrial Applications

The 2001 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Industrial Applications (Preconditioning 2001) was the second conference of its kind to focus on preconditioning techniques for solving various sparse matrix problems. The first conference, which was viewed by the community as a big success, was held at the University of Minnesota, Minneapolis, in June 1999.

The Preconditioning 2001 Conference was held at the Granlibakken Conference Center in Tahoe City, California, on April 29 - May 1, 2001. About seventy participants attended the conference; many were leading experts in the area of preconditioning techniques. The conference featured nine invited plenary presentations, twenty-three contributed talks, and fourteen poster presentations. Overall, the presentations were of high quality, and there was much interaction among the conference attendees.

The attendees came from academia, research laboratories, and industries. The U.S. Department of Energy (DOE) was well represented at the conference; there were participants from six of the major DOE research laboratories (Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge National Laboratory, and Sandia). Some participants were from research laboratories overseas, including Japan, France, and the United Kingdom. Participants from industry included Bell Labs, Boeing, Chevron, Fujitsu, Schlumberger, and Xerox PARC. Academic participants came from universities in the United States, Belgium, France, Germany, the Netherlands, Spain, Switzerland, Sweden, and the United Kingdom.

The conference received support from DOE Headquarters, Lawrence Berkeley, Lawrence Livermore, Sandia, and the University of Waterloo. The conference was also sponsored by the SIAM Activity Group on Linear Algebra.

It is well known that the kernel of many large numerical simulations is the solution of large sparse matrix problems, including in particular the solution of systems of linear equations and the solution of eigenvalue problems. It is not uncommon for the solution of matrix problems to dominate the cost of the entire simulation. Efficient solution of matrix problems therefore has become extremely crucial and will make a big impact on large-scale scientific and engineering simulations.

Examples of such simulations include petroleum refinery processes, environmental modeling (e.g., global warming, groundwater contaminations), structural dynamics, finance, design of advanced semiconductor devices and circuits, accelerator design, fusion energy systems, and aerodynamics.

The traditional means of solving sparse matrix problems is direct factorization. For matrices of small to medium sizes, direct approaches are often adequate. However, as demand for high resolution in the simulations increases, the size of the matrix problems also increases. For extremely large sparse matrices, direct approaches are no longer viable as the amount of memory required can vastly exceed the memory that is otherwise required to represent the application, particularly for three-dimensional simulations. Even if

Synopsis of Workshop Events (continued):

enough memory were not the bottleneck, the time required to solve such large matrix problems would become prohibitive.

An alternative is to use iterative methods to solve these extremely large sparse matrix problems, and research effort has been poured into powerful and efficient iterative methods for decades. However, known low storage iterative methods alone are not sufficient since they lack robustness or require too many iterations; preconditioning techniques are necessary to improve convergence.

Robust preconditioning techniques is a relatively young field when compared to basic direct methods and iterative methods. Many of the techniques developed during the early years did not have much impact initially because of their simplicity and also because of the relatively small size of the matrices to be solved. However, more and more computational experience indicates that a good preconditioner holds the key to an effective iterative method. The big impact of these simple techniques on the performance of an iterative method has attracted increased attentions in recent years. Parallel computers also generate many new research topics in the study of preconditioning. Many new promising techniques have been reported.

However, the theoretical basis for high performance preconditioners is still not well understood; many existing techniques still suffer from lack of robustness; promising ideas need to be tested in real applications. This was the motivation for holding a conference specifically dedicated to the issues in preconditioning and industrial applications. Researchers and scientists/engineers in this field from academia, industries, and research laboratories were brought together so that they could present the progress, discuss difficult issues, exchange their findings, and explore possible new directions.

The topics covered at the Preconditioning 2001 Conference included:

- Incomplete factorization preconditioners
- Domain decomposition preconditioners
- Approximate inverse preconditioners
- Support graph preconditioners
- Multi-level preconditioners
- Preconditioning in eigenanalysis
- Preconditioning in optimization
- Preconditioning in finite element applications
- Preconditioning in image processing
- Applications in fluid dynamics and magnetohydrodynamics
- Applications in multiphase subsurface flow

The conference program included invited plenary lectures, contributed talks, and poster presentations. Selected papers presented from Preconditioning 2001 will appear following a standard review process in a special issue of Numerical Linear Algebra with Applications.

Program details are located at the conference website at http://www.nersc.gov/conferences/pc2001/

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551